

JUL 15 2004

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE 14.Jul.04	3. REPORT TYPE AND DATES COVERED MAJOR REPORT		
4. TITLE AND SUBTITLE SURFACE AND BURIED LANDMINE SCENE GENERATION AND VALIDATION USING THE DIGITAL IMAGING AND REMOTE SENSING IMAGE GENERATION (DIRSIG) MODEL		5. FUNDING NUMBERS		
6. AUTHOR(S) CAPT PETERSON ERIN D				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) ROCHESTER INSTITUTE OF TECHNOLOGY		8. PERFORMING ORGANIZATION REPORT NUMBER CI04-421		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) THE DEPARTMENT OF THE AIR FORCE AFIT/CIA, BLDG 125 2950 P STREET WPAFB OH 45433		10. SPONSORING/MONITORING AGENCY REPORT NUMBER		
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION AVAILABILITY STATEMENT Unlimited distribution In Accordance With AFI 35-205/AFIT 8-1 DISTRIBUTION STATEMENT A Approved for Public Release Distribution Unlimited		12b. DISTRIBUTION CODE		
13. ABSTRACT (Maximum 200 words)				
20040720 078				
14. SUBJECT TERMS		15. NUMBER OF PAGES 12		
		16. PRICE CODE		
17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT	

Surface and Buried Landmine Scene Generation and Validation using the Digital Imaging and Remote Sensing Image Generation (DIRSIG) Model

Captain Erin D. Peterson^{a,b}, Scott D. Brown^b, Timothy J. Hattenberger^b, John R. Schott^b

^aUnited States Air Force*

^bRochester Institute of Technology, Center for Imaging Science
Digital Imaging and Remote Sensing Laboratory
54 Lomb Memorial Drive, Rochester, NY USA 14623-5604

ABSTRACT

Detection and neutralization of surface-laid and buried landmines has been a slow and dangerous endeavor for military forces and humanitarian organizations throughout the world. In an effort to make the process faster and safer, scientists have begun to exploit the ever-evolving passive electro-optical realm, both from a broadband perspective and a multi or hyperspectral perspective. Carried with this exploitation is the development of mine detection algorithms that take advantage of spectral features exhibited by mine targets, only available in a multi or hyperspectral data set. Difficulty in algorithm development arises from a lack of robust data, which is needed to appropriately test the validity of an algorithm's results. This paper discusses the development of synthetic data using the Digital Imaging and Remote Sensing Image Generation (DIRSIG) model. A synthetic landmine scene has been modeled after data collected on the US Army's Yuma Proving Grounds by the University of Hawaii's Airborne Hyperspectral Imager (AHI). The synthetic data has been created and validated to represent the surrogate minefield thermally, spatially, spectrally, and temporally over the 7.9 to 11.5 micron region using 70 bands of data. Validation of the scene has been accomplished by direct comparison to the AHI truth data using qualitative band to band visual analysis, Rank Order Correlation comparison, Principle Components dimensionality analysis, and an evaluation of the R(x) algorithm's performance. This paper discusses landmine detection phenomenology, describes the steps taken to build the scene, methods utilized to overcome limitations of less than adequate ground truth, and compares the synthetic scene to truth data.

Keywords: DIRSIG, reststrahlen, hyperspectral image simulation, mine detection, long wave infrared

1. INTRODUCTION

Since World War II, the ability of an army to detect a minefield that lay in wait has literally meant the difference between life and death. Not only does this apply to advancing forces during a conflict, but also to the residents of the region after the conflict has ended. This introduces the concept of humanitarian demining, a process to ensure innocent civilians are not seriously injured as a result of forgotten and undetectable minefields. As technology has advanced through the decades, the ability to collect and exploit a wider range of data pertaining to landmines and minefields has advanced as well. This has given scientists new abilities to attempt detection of mines and minefields. With a new arsenal of information at the disposal of scientists, the development of new methods to detect mines has followed. From simple metal detectors to the use of thermal imaging or ground penetrating radar, advances have taken place in the countermining field at a quick pace. The challenge at the forefront of research is testing these various detection techniques to quantify how well each performs or potentially, how well some perform in tandem.

The introduction of novel detection techniques into the countermining community has created a need for test data that provides a safe environment for researchers, without the loss of critical information about a mine or a minefield's "signature". By limiting this discussion of the countermining field to passive electro-optical detection techniques, Synthetic Image Generation (SIG) may prove to be a viable solution to the problem at hand.

*The views expressed in this paper are those of the authors and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

Researchers developing anomaly or target detection algorithms for use on hyperspectral data sets or broadband images, cannot afford to undertake huge experimental efforts to produce data spanning a multitude of imaging conditions for evaluation of algorithm performance. It is simply too cost prohibitive and work intensive. Therefore, the need for accurate, reproducible images or sets of data is paramount. SIG can accomplish this task. Specifically, the Digital Imaging and Remote Sensing (DIRS) group's Image Generation (DIRSIG) model can be used for this purpose. DIRSIG is a first principles based physics model that produces high fidelity radiance images of a scene created by a user.¹ Using DIRSIG to build synthetic images offers the algorithm developer total control and flexibility over the data produced. Therefore, if DIRSIG accurately models the physical interactions between objects within the scene, algorithm developers can use the flexibility of the scene generation process to evaluate algorithm performance in a seemingly endless set of scenarios.

The purpose of this project is to use DIRSIG to produce high-resolution images of a scene that contains surface-laid and buried landmine signatures as well as applicable background objects and clutter. The scene can then be used by algorithm developers as a "training ground" to test algorithm performance. Correctly modeling the physics behind object interactions within the scene is of critical importance. If the underlying principles are correct, a scene with accurate minefield signatures will be created, and mine detection or anomaly detection algorithms will perform equally as well on synthetic and real-world data. Hopefully, by using robust data to train an algorithm or refine an algorithm, a more robust algorithm will emerge.

The work discussed in this paper stems from a U.S. Army Multi-University Research Initiative (MURI) sponsored by the Army Research Organization (ARO). Together with four other universities, RIT hopes to further the understanding of the science behind today's ever-growing landmine detection problem.

The approach described in this effort deals not with the specifics of algorithm development, but the generation of accurate synthetic data. A high-resolution scene has been developed and validated by comparing minefield and background attributes of the synthetic scene to attributes of a known data set. A quantification of the differences between the synthetic to the real has also been accomplished. The goal of this work is to provide a validated scene that accurately represents landmine signatures to an ATR algorithm.

2. BACKGROUND

2.1. Landmine Signatures

A landmine signature refers to the presence of a localized difference between a landmine and its surroundings caused by the landmine itself or the emplacement of the mine. This signature manifests itself differently in the visible or near infrared portion of the electromagnetic spectrum as compared to MWIR or LWIR regions. As this work focuses on the LWIR region for validation, the VIS/NIR region has not been emphasized. Signatures of both surface laid and buried mines are presented as potential detection features utilized by a detection algorithm. As the DIRSIG scene focuses on feeding algorithms, spectral signatures of landmines and their interaction with background objects are critically important. Validation of the scene is focused on this area.

2.1.1. Surface Landmine Signatures

The signature produced by a surface laid mine is directly due to the size, shape, composite material makeup, and thermal properties of the mine. These properties are inherently different than surrounding background objects such as soil, grass, etc. When viewed in the thermal region, these property differences will produce an apparent temperature contrast at the sensor. Algorithm developers use this contrast between the target mine and the background as a detection feature. To effectively model this contrast, detailed information must be known about the physical properties and spectral properties of the target mine as well as background objects. Additionally, potential false alarm targets will have similar, but different physical and spectral properties, producing target-like thermal contrasts. The ability of an algorithm to reduce false alarms depends on its ability to distinguish between subtle differences in thermal contrast, whether it is spectrally, spatially, or a combination of the two.

It is difficult to generally characterize the thermal contrast between a mine and the background, as mines come in many different shapes, sizes and compositions. In addition, background variety can be seemingly endless to include wet soil, dry soil, short grass, tall grass, sand, etc. According to Nivelle and Lhomme², soil type can play an important factor in the development of thermal contrast. They observed a contrast inversion when viewing a surface laid mine on rocky soil compared to the same mine laid on a compost background, all other factors held constant. In the first case, the mine was observed to have a negative contrast and in the latter a positive contrast at that particular time of day.

Another important consideration is the diurnal nature of the thermal contrast. This phenomenon is dependent on incident solar radiation as well as heat transfer due to conduction, convection, and radiation.³ Incident solar radiation contributes to heating the mine and the background at different rates, depending on the thermal inertia of the mine and the background, as well as the emissivity of the surfaces of each. Observed signatures are far from constant, varying with atmospheric conditions. There are two noticeable crossover periods where the thermal contrast between the mine and the background is null. In general, these periods occur just after sunrise and just after sunset, but are heavily influenced by atmospheric conditions.³ Just after sunrise and assuming the mine in question heats or cools faster than the background, solar loading has warmed the mine from a point where the mine is cooler than the background to a point where the temperatures are equal. The mine should continue to heat faster than the background throughout the day, assuming a constant solar load. After sunset, solar loading has ceased. The mine and the background begin to transfer heat to the colder, nighttime sky. Again the mine will cool more rapidly than the background, reaching a point where temperatures equal. From a detection standpoint, these times are not suitable, which presents a significant problem for 24-hour detection capabilities. However, this neutral contrast is only truly valid in one spectral band. Because of emissivity variations between the target and background, contrast may persist in other bands.

2.1.2. Buried Landmine Signatures

A buried mine signature is different than that of a surface-laid mine in that the observable features are not of the mine directly, but rather the impact the mine has on the background. The observed thermal signatures of buried landmines are an apparent temperature contrast between the surface temperature of the soil above the mine and the surface temperature of the soil surrounding the mine. There are two commonly observed thermal effects, namely the surface effect and the volume effect.⁴ The surface effect is associated with the process of disturbing the soil directly above the mine during the emplacement process. Disturbing the soil to emplace a mine causes a change in the density of the soil, such that it will have a lower thermal conductivity when compared to the surrounding undisturbed soil.⁵ The lower thermal conductivity leads to a noticeable thermal contrast between disturbed and undisturbed soil. The surface effect is generally applicable to recently buried mines, as environmental conditions such as rain or wind, will degrade apparent temperature contrast. These effects can last up to three weeks under the right conditions.⁴ In observing the surface effect, a broadband infrared sensor may prove to be the best approach, mimicking the common observation technique for observing apparent temperature contrast of a surface-laid mine and its background.⁵

The volume effect deals with the presence of the thermal mass of a buried mine within the soil. It is observed for as long as the mine is in the soil, but reduced depending on environmental and atmospheric conditions.⁴ The soil volume directly above the buried mine will not heat up or cool down at the same rate as the surrounding soil, due to the influence of the thermal mass of the mine. The effect is more pronounced depending on mine burial depth. The deeper the mine is buried, the smaller the apparent temperature contrast.⁶ This observable temperature contrast at the surface of the soil follows a diurnal cycle similar to the diurnal cycle of a surface laid mine. It has been noted through previous work that the variation in apparent temperature contrast of a buried landmine over the diurnal cycle is lower than that of a surface laid mine.⁵ Obviously the exact differences are influenced by mine type, burial depth, background type and other atmospheric parameters.

The second part in understanding buried landmine signatures relates to the spectral structures of the soil disturbed during the burial process and the undisturbed surrounding soil. Immediately after burial of a mine, the disturbed soil will exhibit a localized texture difference capable of being observed by a broadband IR sensor or even in the visible portion of the spectrum.⁷ However, detection based on this feature alone can result in false alarms due to naturally occurring texture differences. Additionally, changing environmental conditions such as heavy rainfall or blowing wind will effectively eliminate any observable texture difference due to mine burial. The main theory behind buried landmine detection due to spectral properties of soil depends on the differences in spectral structure between the surface layer of soil and the subsurface soil.⁷ Essentially, during the burial process the subsurface soil is churned, resulting in some of the subsurface soil residing at the surface. A spectral difference can be observed between subsurface soil and surface soil based on weathering effects on the surface soil and organic composition differences. Over the MWIR and LWIR regions, soil will show spectral structure due to specific spectral features of the minerals contained within.⁸ Therefore, the simplest method of detection would be to observe a different spectral signature from the disturbed soil based solely on a change in mineral composition.

If the disturbed and undisturbed soils are of equal mineral content, all hope is not lost. A spectral feature common to most soils is the silicate reststrahlen feature, which manifests itself in the 8.5 to 9.5 micron spectral window.⁸ This feature can be exploited to detect buried objects. Soil particle size plays an important role in determining an observable emissivity difference in the reststrahlen bands between disturbed and undisturbed soil. Before soil

disturbance occurs, the subsurface soil layer is composed of a mixture of larger and smaller particles. The surface layer, having been exposed to wind, rain, and other atmospheric effects, is composed of only large particles. It has been shown that a higher emissivity value is observed for the subsurface soil composition, i.e. the disturbed soil, in the reststrahlen bands compared to undisturbed soil.⁷ The spectral difference between untouched soil and soil that has been disturbed during the landmine burial process is the key feature exploited by spectrally based mine detection algorithms.

3. SCENE MODELING

To synthetically represent a minefield with the spectral and spatial complexity of field collected data, a robust set of fundamental object and material parameters must be incorporated. If not, the simulation will fail to adequately represent real world image data, ultimately leading to unrealistic algorithm performance. Hence the modeler's creed: garbage in equals garbage out. Unlike previous modeling work⁹ dealing with complex object and background interactions in the reflective region of the spectrum, this paper addresses the unique and somewhat more complex issues encountered while modeling in the LWIR region of the spectrum. This paper also addresses novel modeling techniques that have been incorporated to overcome limitations of less than adequate ground truth, a prevalent occurrence throughout the creation of this scene. With that said, the goal of the project was not to exactly replicate the truth data, but to develop solutions to complex modeling challenges as well as provide statistically realistic data to mine detection algorithms with the purpose of furthering algorithm development. Exact one-to-one scene correlation was not required or expected. Four versions of the scene have been rendered, one from an altitude of 700 feet AGL during the evening (approximately 1700hrs), one from an altitude of 1400 feet AGL during the evening, one from an altitude of 700 feet during the middle of the day (approximately 1300hrs, referred to as 'noon'), and one from an altitude of 1400 feet during the middle of the day. These renderings coincide with data collections over a simulated minefield at the US Army's Yuma Proving Grounds in Arizona, USA by the University of Hawaii's Airborne Hyperspectral Imager (AHI). The AHI data consisted of calibrated radiance images covering the LWIR from approximately 7.5 to 11.5 μ m, utilizing 70 spectral bands. As will be discussed in further detail, this data was invaluable during the scene creation process.

3.1. Scene Geometry

From the available ground truth, the scene's background objects and target spatial dimensions were estimated then created in facetized model form using a computer aided design (CAD) software package (see figures 1-3). Six (6) objects in all were built and each placed in the scene corresponding to surveyed target location coordinates. These included simulant M19 landmines, simulant M20 landmines, Top Hat fiducial markers, EO/IR calibration panels, large desert bushes and small desert bushes.




Figure 1. A facetized M19 landmine model.




Figure 2. A facetized M20 landmine model.




Figure 3. A facetized desert bush model.

In order to adequately represent small-scale surface elevation features (bumps in the ground), a new technique was employed, as standard scene creation methods would not provide the necessary degree of resolution. For this scene we used a bump-map, which is based off standard mapping routines¹⁰ but used to alter the direction of the surface normal vector of a facetized object (see Figure 4). The amount and direction of surface normal deflection is based upon the ground sample distance and the gradient of gray level values across the mapping image used in the process. The steeper the gradient, the higher the angle of deflection applied to the surface normal. This produces the appearance of a roughened or bumpy surface without the need to explicitly build surface variation into the CAD model. In this case, a bump-map was applied to soil areas in the scene to give it a natural and realistic look. This process also affects the amount of solar loading onto the object, where the surface variation will incorporate changing angular effects across the surface of an object. This produces changes in temperature across an object's surface and leads to temperature variation across the scene.




Figure 4. An example of surface normal deflection to produce bumps in a flat surface.




Figure 5. A buried mine area's material map.

To facilitate the creation of buried mine areas, a material map was introduced that was sufficiently large enough to completely cover a 150 by 150 meter region using a ground sample distance of 0.025 meters. Within the map, small regions were inserted corresponding to the exact spatial location of each surveyed area, using two (2) gray levels to represent disturbed soil (see Figure 5). Gray regions represent disturbed soil that is located directly above the buried mine and black regions represent surrounding disturbed soil. Any additional white pixels represent undisturbed soil.

3.2. Sensor Modeling

In addition to standard sensor modeling within DIRSIG¹, a unique approach was taken to include realistic spectrally correlated noise derived from truth imagery into the final scene renderings. Limited information pertaining to the noise characteristics of the AHI sensor was known, so this method of creating realistic noise was the only way to add this variability into the scene. The process is in two parts. Step 1 derives an estimated dark scan from truth data and step 2 uses the statistics from the estimated dark scan to produce synthetic noise cubes that can be added on a band by band basis to the rendered DIRSIG imagery. The flow for this process is presented in Figures 6 and 7.

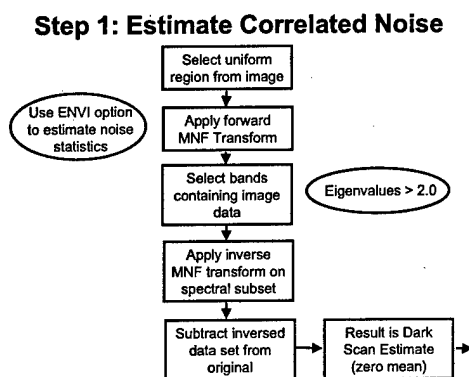


Figure 6. The noise development process - Step 1.

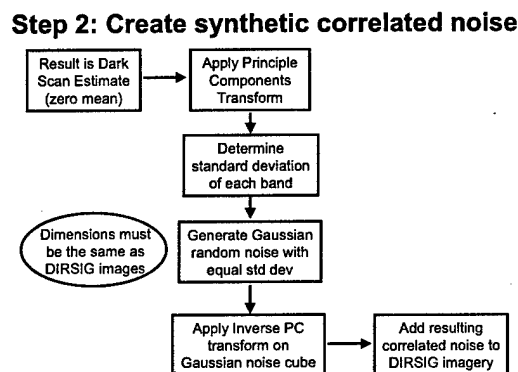


Figure 7. The noise development process - Step 2.

Initially, a materially uniform region is selected for processing from the truth data. Next, a standard Minimum Noise Fraction (MNF) data transformation is employed using ENVI¹¹. ENVI has a built-in function to estimate noise statistics from the data at hand through a shift-difference method. This method differences neighboring pixels directly above and to the right of the base pixel, then averages the values to obtain the base pixel's noise value. This process derives an estimate of scene noise for use in the MNF transform. The resulting "image" bands of data are segmented from the "noise" bands by selecting only the bands with eigenvalues greater than 2.0, as eigenvalues near 1.0 indicate pure noise. The noiseless bands are inversely transformed back to the image data space and then subtracted from the original uniform region data to obtain an estimated dark scan for the imaging system. The resulting noise data is inherently spectrally correlated with a mean value of zero. This noise cannot be directly added to rendered DIRSIG imagery because the noise image does not necessarily match the rendering's exact image dimensions. To rectify this and to retain the inherent spectral correlation, the Step 2 process is employed. A standard Principle Components (PC) transformation is applied to the estimated dark scan, which de-correlates the data. The standard deviation for each transformed band of data is determined and then used to generate synthetic Gaussian distributed noise images having equivalent image dimensions as the DIRSIG rendering. An inverse PC transformation is applied to the synthetic noise cube, using the transformation statistics from the estimated dark scan's forward transform. This step correlates the synthetic noise such that it is nearly identical to the input correlation. The synthetic, correlated noise is then directly added to the DIRSIG rendering on a band-by-band basis. Pictorial examples of the resulting covariance matrix compared to the input covariance matrix of the estimated dark scan are presented in figures 8 and 9.




Figure 8. (A) E700 AHI derived noise covariance, (B) E700 synthetic noise covariance.




Figure 9. (A) N700 AHI derived noise covariance, (B) N700 synthetic noise covariance.

3.3. Thermal material parameters

The DIRSIG material file contains all the needed thermodynamic information about a single material for DIRSIG's thermal model to develop a diurnal temperature prediction. These parameters include fundamental material properties such as specific heat, thermal conductivity, mass density, thickness, and broadband emissivity values for the visible portion and thermal portion of the spectrum, as well as modeling parameters such as specularly or exposed area. The basic parameters for all materials within the scene were initially estimated from previously validated data at Rochester Institute of Technology or from textbook material parameters, due to a lack of supporting ground truth. Refinements to the initial values were made throughout the scene creation process in order to more closely align final synthetic radiance curves to truth radiance curves.

Adequately representing thermal clutter and thermal variation within a material played a significant role in developing the synthetic scene. In order to ensure anomaly or target detection algorithms perform consistently between truth and synthetic data, the issue of thermal variation had to be addressed. Previous hyperspectral modeling with focus on algorithmic development⁹ addressed issues encountered in the visible portion of the spectrum. This current work is one of the first attempts to address the issues of thermal clutter and thermal variation conducted by the DIRS group.

After initial evaluation of the scene using the baseline material mapping methodology, it was determined that additional areas of compacted soil were needed to add additional thermal variation and clutter across the landscape. To add these areas of soil, two additional gray level values were added to the base material map to represent medium-compact and highly compact soil (see figure 10). This facilitated some thermal variability, but not to the degree needed. To rectify the situation, the baseline undisturbed soil (all white space in figure 10) was split into three variants. The first variant was the baseline undisturbed soil and the second was a "low temperature" version created by reducing the broadband visible emissivity value of the baseline by 0.03 and increasing the broadband thermal emissivity value by 0.03. The third variant was a "high temperature" version created by increasing the broadband visible emissivity value of the baseline by 0.03 and decreasing the broadband thermal emissivity value by 0.03. In order to incorporate these new materials, the material map was again altered. A three level threshold was applied to a variable gray level image, creating equal amounts of each undisturbed soil variant. The resulting image (see figure 11) was used to replace the baseline undisturbed soil's gray level value in the base material map.




Figure 10. The DIRSIG material map with gray areas of added thermal clutter. Note the rows of buried mines.




Figure 11. A variable image with a three-level threshold applied.

By adding new materials to the scene, additional thermal clutter and variation were incorporated. Prior to the creation of these materials, the synthetic scene showed thermal variation of approximately 0.1°C in the evening, 700-foot rendering and approximately 2.5°C in the noon, 700-foot rendering, due only to surface geometric effects introduced through bump mapping. After the incorporation of the new material map, thermal variation in the evening, 700-foot rendering was approximately 0.3°C and approximately 3.5°C in the noon, 700-foot rendering. This compares to thermal variation in the truth data of approximately 0.4°C in the evening, 700-foot data and approximately 3.0°C in the noon, 700-foot data. This represents a significant improvement in the thermal modeling process.

3.4. Optical material parameters

Valid emissivity curves for each material within the scene are absolutely essential for successful reproduction of observed material features. In addition, key buried mine detection features reside primarily between differences in emissivity curves, so valid curves become even more important. Unfortunately, measured emissivity curves for materials within the scene were unavailable, leaving scene builders with quite the dilemma. To rectify this situation, two options existed. The first to develop or research a method to extract emissivity curves from the truth data or the second to utilize emissivity curves of similar, but not equivalent, materials. To facilitate better comparison between material radiance curves and the evaluation of the modeling process, the former method was selected. Utilizing emissivity curves gathered from other sources may have been sufficient for man-made materials such as surface mines, fiducial markers, or calibration panels, however soil composition would have been extremely difficult to match. The selection of any externally gathered emissivity curve would have been more of a guess than a scientific selection process. To ensure a sense of uniformity in selecting emissivity curves, the decision was made to use a single technique for all materials rather than a mix of two or more.

Deriving emissivity curves from the truth data was a unique process in itself due to the spectral region being modeled. The process entails determining atmospheric components of the sensor reaching radiance then applying a blackbody curve fit technique to radiance curves that have been atmospherically calibrated. To determine values for the average atmospheric transmission and average upwelled radiance across the scene, a multiple altitude calibration method¹² was employed. The limited amount of truth data was nearly perfect for the use of this technique. Two of the AHI calibrated radiance images were used, specifically both evening shots at altitudes of 700 and 1400 feet respectively. These data sets were obtained over the course of approximately 14 minutes and covered the same general area on the ground. Given the low nature of the flights and small sensor field of view, angular correction for off-nadir imaging was assumed to be unneeded. In addition to the reasons stated previously, these sets of data were chosen primarily because target pixels were discernable in both sets of data and these data sets were imaged at a time of day that had low thermal variability, with no issues arising from solar loading, since the sun had already set. Most observed radiance variability from a material at this time of day was due to inherent differences in emissivity. The result of this process produced average scene atmospheric transmission and upwelled radiance values across the spectral region sensed by the AHI system, shown in figure 12. Pixels corresponding to specific materials were located in the truth imagery and interrogated for their radiance curve. The radiance curve was atmospherically corrected using the previously discussed values, and then became the input for the blackbody curve fit procedure, namely the Planck Curve Fit¹².




Figure 12. Average atmospheric transmission and upwelled radiance curves across the AHI truth scene at 700 feet.




Figure 13. Average derived emissivity curves for disturbed and undisturbed soil.

The Planck Curve Fit is a temperature-emissivity separation technique, which involves fitting a target's radiance curve to Planck's blackbody radiation curve generated at the highest temperature that will keep the target's emissivity from

exceeding unity. This is an iterative process where two initial high and low temperature bounds are selected and then a blackbody radiance curve corresponding to the mid-point of the temperature range is generated and compared to the target radiance curve. This midpoint temperature then replaces either the upper or lower temperature bound and the process repeats. Iterations continue until the upper and lower bounds are separated by a defined value, e.g. 0.5°K or similar range. This method is quick and easy to implement, but hinges on at least one point along the target's emissivity curve approaching unity. For the case of soil emissivity over the LWIR spectral region this assumption holds true, however this does not hold true for manmade materials such as the fiducial markers in the scene. Lacking alternative data, this process was implemented to derive families of emissivity curves for all discernable materials within the scene. Families of curves meaning a number of derived curves that encompass the range of spectral variation for a given material. As an example, twenty curves were extracted to represent disturbed soil as a whole. The results of the process can be seen in figure 13.

The application of emissivity curves to material types was accomplished through the use of a texture map¹⁰, by which emissivity variation within a material was introduced into the scene. Determining a good gray level image as the basis for the application of texture was a difficult process. When texturing in the reflective region, one only needs to worry about emissivity differences when attempting to spatially match observed brightness variability across a material. Typically, an overhead grayscale image of the true scene is used to drive the texture process, ultimately determining how well a rendered image looks. In the LWIR, the previous issue is complicated by temperature variability as well. In this work, matters are complicated even further because there was no available overhead imagery that was suitable for texturing. Therefore, any input texture image would be nothing more than an educated guess at what the true spatial distribution of material emissivity curves truly is. A generic texture image was selected that attempted to match the spatial distribution of emissivity values observed in the evening truth imagery. For the same reason that the evening data was selected to derive the emissivity curves, they were also chosen for attempting to match texture images. The thought being that radiance variability is driven predominantly by emissivity differences due to low temperature variability at that particular time of day. The problem of texture application in the thermal infrared is an area that is quite ripe for further research. Hopefully this modeling effort can provide a starting point for future work in LWIR texturing.

4. RESULTS

As previously mentioned, four data sets were created for qualitative and quantitative evaluation to truth data. It would be impossible to present the details of every evaluation, so the following results address the most relevant points. The metrics used to evaluate the scene range from the purely qualitative to purely quantitative. Both are necessary for a thorough discussion of the pros and cons of the current scene creation process.

4.1. Visual image comparison

The first approach to evaluating the goodness of fit between synthetic and truth data is to simply view the rendering and truth imagery side by side. A simplistic qualitative comparison is the best method available for noticing major global differences between the two. Because the data is hyperspectral, only two of the 70 available bands are shown in the comparison, namely bands 25 and 65. These bands lie at approximately 9.16 μ m and 11.22 μ m respectively, the first within the reststrahlen band feature and the second outside of the feature. These are presented in figures 14–17, where figures 14 and 15 address low altitude, evening versions of the scene, and figures 16 and 17 address midday imaging conditions. Evening data shows appropriate levels of contrast between target areas (brighter areas in vertical rows) and background areas combined with realistic levels of clutter. In addition, target areas diminish when viewing the scene off the reststrahlen band feature, as expected. Daytime data shows the difficulty in adequately matching spatial temperature variation throughout the scene. Especially noticeable in figure 17, background brightness variation is spatially on a wider scale in the rendering, due to the spatial distribution in the material map. As discussed previously, actual variation in temperature values is appropriate between the truth and synthetic data, however the actual spatial distribution of these values is a difficult process to perfect. This is the overarching problem encountered with adding temperature variation across a material in the manner previously described. Also, in figures 14 and 15, the time of day the imagery was collected or rendered allows for a good comparison of the emissivity texturing process. Again, without overhead imagery of the scene the rendering will not exactly match the truth, nor was this the goal. However, the generic texture image must be reasonable to produce a viable scene. Since temperature variation is at a minimum at this time, radiance variation is due mostly to emissivity differences. This qualitative comparison concludes that the general texture image approximates the truth imagery quite sufficiently and is a viable modeling option. Additionally, the structured noise, or vertical striping as seen in the AHI data, was not modeled in the synthetic data. Any attempt to characterize this band

specific and spatially varying noise would require detailed knowledge of the sensor's focal plane at the time of imaging, knowledge that was unavailable for this work.




Figure 14. Both images are evening at 700 ft., (L) Band 25, AHI truth data image; (R) Band 25, DIRSIG rendering.




Figure 15. Both images are evening at 700 ft., (L) Band 65, AHI truth data image; (R) Band 25, DIRSIG rendering.




Figure 16. Both images are noon at 700 ft., (L) Band 65, AHI truth data image; (R) Band 65, DIRSIG rendering.




Figure 17. Both images are noon at 700 ft., (L) Band 65, AHI truth data image; (R) Band 65, DIRSIG rendering.

4.2. Rank order correlation

One of the most effective evaluations conducted was an evaluation of rank order correlation¹⁴ amongst materials within the scene over each band of data. The scene is to be processed by algorithms, keying off likely contrast features in the minefield created by either a thermal difference or emissivity difference. This difference in brightness value is the quantity that is of most importance to correctly model. An attempt to strictly quantify the mean level error between truth and synthetic data will not provide an accurate assessment of how well the synthetic data approximates the truth to an algorithm. As an example, if RMS error were a measure used for evaluation, a constant gain or bias error between truth and synthetic data would yield poor results. If the contrast between scene materials were the same, any global gain or bias would be insignificant to an algorithm. Rank order correlation provides an avenue to evaluate the in-scene contrast while removing error contributions from less significant sources. In order to evaluate the relative brightness of each scene compared to truth, six materials were selected for comparison. In each band of each scene, these six materials were ranked brightest to darkest, then evaluated using the Spearman rank order correlation¹⁴. The metric values range from -1.0 to 1.0 with 1.0 being perfect correlation, the optimum result for this evaluation. Each band of data carries a unique correlation value, all of which were averaged to produce a single value representative of the entire hyperspectral scene. These results are shown in Table 1. The results for this portion of the evaluation indicate that brightness contrast between materials can be faithfully reproduced using the DIRSIG model. Minimum values listed for the evening renderings are a singular occurrence, observed in the data's first spectral band, lying on the edge of a water vapor absorption feature. The low values are primarily due to the Planck curve fit procedure, as each material's emissivity approaches unity in the initial spectral band. At this time of day where thermal differences are at a low, average radiance

values were very tightly grouped, resulting in a relatively arbitrary brightness ranking. Minimum values in the midday data also arise due to a very tight grouping of radiance values for three different scene materials. Minimum values tend to predict issues with the emissivity derivation process or represent situations where the metric struggles to appropriately rank materials. In addition to overall scene values, each material was evaluated, ranking the material's radiance value in each spectral band. A rank order correlation value was determined for each material, in each scene and finally averaged for an overall material "score". This evaluation tends to measure spectral correlation determined by atmospheric constituents and emissivity curves. Essentially, this metric evaluates how well the atmosphere matches the truth atmosphere and how well the derived emissivity curves match the truth. Results for this evaluation are presented in Table 2. With the exception of the Top Hat fiducial markers, all individual DIRSIG materials show strong correlation with the truth. Strong correlation across all but one material implies that the atmosphere is modeled quite well. It also points to the Top Hat's derived emissivity curve as the flaw in the process. This was not unexpected, since the Planck curve fit used to separate temperature from emissivity is predicated on the target's emissivity approaching unity, which is typically not the case for man-made materials such as the Top Hat. Taking this one step further, the data implies that the emissivity derivation process tends to break down for non-Lambertian materials. Soils and vegetation, being essentially Lambertian, have distinctly high correlation values. The Top Hat fiducials and EO/IR panels are not as Lambertian. Particularly, the Top Hat fiducials are quite specular, being made from a polished metal substance.

Overall Scene Rank Order Correlation (band average)				
	Mean	Max	Min	Standard Deviation
E700	0.93	1.00	0.43	0.09
E1400	0.94	1.00	0.40	0.09
N700	0.86	1.00	0.60	0.09
N1400	0.93	1.00	0.60	0.10

Table 1. Overall scene rank order correlation statistics, comparing DIRSIG imagery to truth imagery.

Material Rank Order Correlation					
	Undisturbed Soil	Disturbed Soil	EO/IR Panel	Vegetation	Top Hat Fiducial
E700	0.96	0.98	0.91	0.94	0.65
E1400	0.95	0.97	0.87	0.92	0.77
N700	0.85	0.96	0.85	0.96	0.37
N1400	0.86	0.91	0.82	0.94	0.43
AVG Value	0.91	0.95	0.86	0.94	0.55

Table 2. Individual material rank order correlation values, comparing DIRSIG imagery to truth imagery.

4.3. Dimensionality analysis

One of the simplest, but most important analyses is an evaluation of dimensionality. The project goal is a good overall statistical correlation between synthetic and truth data. The amount of inherent variability in the data directly compares to the statistical fit between the two. To evaluate this, a standard Principle Components transformation was applied to each of the truth and rendered data sets. The Principle Components transformation¹² is used to decorrelate and maximize data variability in multi or hyperspectral data sets. This is accomplished by projecting the original multi-band data onto a new set of orthogonal axes defined by the eigenvectors of the data set. The resulting PC data bands will be ordered such that the first PC band contains the largest percentage of data variance, the second PC band contains the second largest percentage of data variance, and so on. If the rendering is a good fit to the truth, the resulting amount of variance in each PC band should be similar between the real and synthetic data. If the synthetic captures most of the data variance in one or two bands, where the truth spreads the variance over six to seven bands, we can conclude that the synthetic is far less statistically complex and will not approximate the scene appropriately to an algorithm. The results of the analysis are shown in Table 3, and are very encouraging. The evening data compare extremely well, even showing slightly more spread in the synthetic variability. The DIRSIG midday data is not as variable as one would hope. This seems to be due to issues arising with the addition of solar loading in the scene. As mentioned previously, the sun had set at the time of the evening renderings, eliminating solar issues from the data. In the midday data, the added complexity of the sun adds to the thermal variability in the soils. The solution to lacking thermal variability was to increase the number of slightly different soil materials. While this did a reasonable job of approximating the thermal variability, the complexity of the issue may not have been fully replicated. These results point to the need for additional flexibility in the DIRSIG model when adding temperature variation, an issue currently under investigation. Additionally, some of the reduction in variability in the noontime data may result from complex surface variation not fully captured by the scene's bump map. In terms of scene building, this could be eliminated with a detailed knowledge of the elevation changes across the scene, incorporated into a Digital Elevation Map or DEM. A DEM combined with a

bump map would produce additional surface variability. In this case there was no prior knowledge of the surface variation across the truth scene, so a DEM was unavailable.

PC Band	Percentage of Variance							
	E700		E1400		N700		N1400	
	AHI	DIRSIG	AHI	DIRSIG	AHI	DIRSIG	AHI	DIRSIG
1	85.80%	79.74%	76.28%	72.99%	84.72%	93.86%	84.53%	96.11%
2	8.64%	14.32%	13.83%	15.75%	13.36%	5.15%	13.38%	3.11%
3	1.12%	1.45%	1.92%	1.28%	0.71%	0.44%	0.66%	0.25%
4	0.53%	0.60%	0.66%	0.75%	0.29%	0.21%	0.35%	0.12%
5	0.33%	0.22%	0.55%	0.68%	0.23%	0.02%	0.22%	0.02%
6	0.16%	0.21%	0.27%	0.64%	0.07%	0.02%	0.06%	0.02%
7	0.14%	0.19%	0.24%	0.49%	0.04%	0.01%	0.05%	0.02%
8	0.12%	0.17%	0.23%	0.46%	0.03%	0.01%	0.03%	0.02%
9	0.12%	0.17%	0.23%	0.44%	0.03%	0.01%	0.03%	0.02%
10	0.11%	0.17%	0.22%	0.41%	0.02%	0.01%	0.02%	0.02%

Table 3. A comparison of data variance between the AHI truth data and the DIRSIG synthetic data across all four data sets.

4.4. Comparative algorithmic performance

The final evaluation is a comparison between an anomaly detection algorithm's performance on the truth scene and the synthetic scene. Many "canned" anomaly detection algorithms exist today, but in an attempt to provide a more strenuous test, the R(x) anomaly detection algorithm¹⁵ was decided upon. This algorithm processes a multi or hyperspectral data set spectrally and spatially over a user defined processing window size. This is unique in that standard anomaly detection algorithms tend to process either spectrally or spatially, but not in a combined fashion. It is important to note that this algorithm is not designed specifically to address the mine detection issue; so excellent mine detection was not a requirement. The goal was to see if the algorithm would perform equally as well on both data sets, whether that be good or bad. A generic processing window size of 21x21 pixels was used for all evaluations, with a target spatial shape defined as a 5x5 square set of pixels for the 700 foot data and a 3x3 square set of pixels for the 1400 foot data. These dimensions were selected to tune the algorithm to the approximate spatial size of the buried mine areas in the truth imagery. Comparison images of the midday, 700 foot results and the evening, 1400 foot results are presented in figures 18 and 19. A 98% image data threshold has been applied, allowing a visual comparison.




Figure 18. R(x) algorithm result images for (L) AHI truth data and (R) DIRSIG rendering (without added roll or roll correction) for midday, 700 ft.




Figure 19. R(x) algorithm result images for (L) AHI truth data and (R) DIRSIG rendering (without added roll or roll correction) for evening, 1400 ft.

The result images are very encouraging. The number of false alarms is very significant in both the truth and synthetic data. While the results are not identical, they show that spectral and spatial clutter has been created that can adequately model the real world. The amount of clutter seen in the synthetic data, while appearing to be less than the truth, is of lower concern from a modeling standpoint because additional clutter objects can easily be inserted. The fact that these clutter objects approximate truth clutter quite well spatially and spectrally is the significant point. It was expected that

mined areas would be more noticeable in the synthetic data, purely due to the fact that it is modeled data, however some of this can be explained. Buried mine areas were created to have a very similar circular shape, where true buried mine areas are quite irregular. Adding more irregularity to these areas would have a dramatic effect in the algorithms ability to detect. As mentioned before, the target's spatial structure was input to the algorithm as a regular square of pixels, which will more closely match the synthetic data than real world. From a scene building prospective, this scene more closely represents the truth to a target detection algorithm than any other thermal scene built at the DIRS lab to date. That being said, this scene is only the first step to exact statistical representation of the truth and will be improved upon.

5. CONCLUSION

This work has demonstrated that DIRSIG is capable of accurately modeling a representative LWIR scene complete with spectrally and spatially varying clutter sources, even when lacking significant amounts of truth data. New techniques have been developed and used to assist the scene builder in deriving accurate input data for the model, such as synthetic correlated noise generation and emissivity curve derivation. Moreover, this work has shown the need for further study in the area of temperature variability modeling combined with emissivity texturing. Scene modeling is cyclical in nature; a first version is created and evaluated, and then upgrades are determined and implemented. The new scene is rendered and the process repeats. With each cycle, the scene gets closer and closer to the truth. While this may never be perfectly attainable, striving to understand the error takes us a few steps closer to equality. The next phase of effort for this work will be to investigate methods to incorporate suitable temperature variability as well as to vary the spatial dimensions and locations of buried mine areas within the scene. The focus will be improving the scene so as to enhance the development of mine detection algorithms.

REFERENCES

1. Schott, J.R., Brown, S.D., Raqueño, R.V., Gross, H.N., and Robinson, G., "An advanced synthetic image generation model and its application to multi/hyperspectral algorithm development", *Canadian Journal of Remote Sensing*, Vol. 25(2), pp. 99-111, 1999.
2. Nivellet, F. and Lhomme, P., "Detection of land mines with passive IR and mmW imaging sensors", *Proceedings of the SPIE*, Vol. 3079, pp. 614-624, 1997.
3. Maksymonko, G., Ware, B., Poole D., "A characterization of diurnal and environmental effects on mines and the factors influencing the performance of mine detection ATR algorithms", *Proceedings of the SPIE*, Vol. 2496, pp. 140-151, 1995.
4. Simard, J.R., "Improved landmine detection capability (ILDC): Systematic approach to the detection of buried mines using passive IR imaging", *Proceedings of the SPIE*, Vol. 2765, pp. 489-500, 1996.
5. DePersia, A., Bowman, A., Lucey, P., Winter, E., "Phenomenology considerations for hyperspectral mine detection", *Proceedings of the SPIE*, Vol. 2496, pp. 159-167, 1995.
6. Khanafer, K. and Vafai, K., "Thermal analysis of buried land mines over a diurnal cycle", *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 40(2), pp. 461-473, 2002.
7. DePersia, A., et al., "ARPA's Hyperspectral mine detection program", *Proceedings of the Third International Symposium on Spectral Science Research (ISSSR)*, 1995.
8. Winter, E., et al., "Experiments to support the development of techniques for hyperspectral mine detection", *Proceedings of the SPIE*, Vol. 2759, pp. 139-148, 1996.
9. Ientilucci, E.J., and Brown, S.D., "Advances in wide area hyperspectral image simulation", *Proceedings of the SPIE*, Vol. 5075, pp. 110-121, 2003.
10. Brown, S.D., and Schott, J.R., "Characterization techniques for incorporating backgrounds into DIRSIG", *Proceedings of the SPIE*, Vol. 4029, pp. 205-216, 2000.
11. "The Environment for Visualizing Images (ENVI)" software, information online at <http://www.rsinc.com/envi/>.
12. Schott, J.R., *Remote Sensing: The Image Chain Approach*, pp. 198-200, Oxford University Press, NY, 1997.
13. Kahle, A.B. and Alley, R.E., "Separation of temperature and emittance in remotely sensed radiance measurements", *Remote Sensing of the Environment*, Vol. 42, pp. 107-111, 1992.
14. Mason, J., Schott, J., Salvaggio, C., Sirianni, J., "Validation of contrast and phenomenology in the Digital Imaging and Remote Sensing (DIRS) lab's Image Generation (DIRSIG) model", *Proceedings of the SPIE*, Vol. 2269, pp. 622-633, 1994.
15. Reed, I.S. and Yu, X., "Adaptive multiple-band CFAR detection of an optical pattern with unknown spectral distribution", *IEEE Transactions on Acoustics, Speech, and Signal Processing*, Vol. 38(10), pp. 1760-1770, 1990.